

CLAIMS

1. A method for use in a computer system including a plurality of host computers including a root host computer and at least one child host computer, the root host computer having a volume of storage available to it that is stored on at least one non-volatile storage device, the method comprising an act of:

(A) exporting at least a portion of the volume of storage from the root host computer to the at least one child host computer so that the at least one child host computer and the root host computer share access to the volume of storage.

10 2. The method of claim 1, wherein the act (A) includes an act of exporting the at least a portion of the volume of storage to a plurality of child host computers so that the plurality of child host computers and the root host computer share access to the volume of storage.

15 3. The method of claim 1, further comprising an act of:

(B) exporting the at least a portion of the volume of storage from the at least one child host computer to at least one grandchild host computer of the plurality of host computers so that the at least one child host computer, the at least one grandchild host computer and the root host computer share access to the volume of storage, wherein the at least one child host computer has access to the volume of storage through the root host computer, and wherein the at least one grandchild host computer has access to the volume of storage through the at least one child host computer to create a logical hierarchy of host computers that share access to the volume of storage.

25 4. The method of claim 1, wherein the at least one non-volatile storage device is provided on the root host computer.

30 5. The method of claim 1, wherein the computer system further comprises at least one storage system that includes the non-volatile storage device and makes the volume of storage available to the root host computer.

6. The method of claim 1, further comprising an act of caching a local copy of the at least a portion of the volume of storage on the child host computer.

7. The method of claim 3, further comprising acts of:
5 caching a first local copy of the at least a portion of the volume of storage on the at least one child host computer; and
caching a second local copy of the at least a portion of the volume of storage on the at least one grandchild host computer.

10 8. The method of claim 3, wherein the at least one child host computer comprises a plurality of child host computers, and wherein at least one of the plurality of child host computers is untrusted by the root host computer.

15 9. The method of claim 3, wherein the at least one child host computer has access to the volume of storage through the root host computer and the at least one grandchild host computer has access to the volume of storage through the at least one child host computer to create a logical hierarchy of host computers that share access to the shared volume, the logical hierarchy have at least a root level, a child level and a grandchild level, and wherein the method further includes an act of:

20 (C) exporting the at least a portion of the volume of storage to at least one of the plurality of host computers in the logical hierarchy that is untrusted by the root host computer.

25 10. The method of claim 6, wherein the act of caching a local copy of the at least a portion of the volume of storage on the child host computer includes caching the at least a portion of the volume of storage in a set associative cache having at least two separately indexed portions.

30 11. The method of claim 5, wherein the volume of storage is a logical volume made available from the at least one storage system to the root host computer.

12. The method of claim 11, wherein the at least one non-volatile storage device is a disk drive in the at least one storage system.

13. The method of claim 5, wherein the volume of storage is a logical volume
5 made available from the storage system to the root host computer, and wherein the at least a portion of the volume of storage is a block of the logical volume.

14. The method of claim 3, wherein the computer system further comprises a network that interconnects the root host computer, the at least one child host computer
10 and the at least one grandchild host computer.

15. The method of claim 3, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, and wherein the method
15 further comprises an act of:

storing a local copy of at least portions of the volume of storage for each of the plurality of child host computers and the plurality of grandchild host computers.

16. The method of claim 15, further comprising an act of:
20 in response to a write request from one of the plurality of child host computers and the plurality of grandchild host computers to write to a first portion of the volume of storage, invalidating any local copies of the first portion of the volume of storage associated with the other child and grandchild host computers prior to authorizing the write request to proceed.

25
17. The method of claim 3, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy
30 for distributing the shared volume of storage throughout the computer system, and wherein the method further comprises an act of:

providing each node in the logical hierarchy with information that enables it to participate in the hierarchy, wherein the information is limited to information relating only to adjacent nodes in the hierarchy.

5 18. -- The method of claim 1, further comprising an act of storing a local copy of the at least a portion of the volume of storage for the at least one child host computer.

19. The method of claim 3, further comprising acts of:
 storing a first local copy of the at least a portion of the volume of storage
10 for the at least one child host computer; and
 storing a second local copy of the at least a portion of the volume of storage for the at least one grandchild host computer.

20. The method of claim 3, wherein the computer system has a physical configuration, and wherein the logical hierarchy is independent of the physical configuration of the computer system.

21. The method of claim 3, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy for distributing the shared volume of storage throughout the computer system, and wherein the method further comprises acts of:

 storing a local copy of the volume of storage for each of the nodes in the hierarchy, each local copy comprising a copy of at least a portion of the volume of storage, wherein the at least a portion of the volume may differ between the nodes in the hierarchy; and

 storing, in association with each local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the associated node to the at least a portion of the volume of storage stored in the local copy.

22. The method of claim 21, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for at least one second node in the hierarchy that is descendant from the at least one first node in the hierarchy.

5

23. The method of claim 21, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for each of the nodes in the hierarchy that is descendant from the at least one first node.

10 24. The method of claim 21, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the method further comprises acts of:

when a failed node in the hierarchy fails, recovering from the failure by reconnecting direct descendant nodes of the failed node to the parent node of the failed
15 node; and

reestablishing the permission information for each of the direct descendant nodes of the failed node from the parent node of the failed node.

25. A method for use in a computer system including a plurality of host
20 computers and at least one storage system, the plurality of host computers including a root host computer and at least one child host computer, the at least one storage system making a volume of storage available to the root host computer, the at least one storage system having at least storage device on which the volume of storage is stored, the method comprising an act of:

25 (A) exporting at least a portion of the volume of storage from the root host computer to the at least one child host computer so that the at least one child host computer and the root host computer share access to the volume of storage.

26. The method of claim 25, wherein the act (A) includes an act of exporting
30 the at least a portion of the volume of storage to a plurality of child host computers so that the plurality of child host computers and the root host computer share access to the volume of storage.

27. The method of claim 25, further comprising an act of:

(B) exporting the at least a portion of the volume of storage from the
at least one child host computer to at least one grandchild host computer of the plurality
5 of host computers so that the at least one child host computer, the at least one grandchild
host computer and the root host computer share access to the volume of storage, wherein
the at least one child host computer has access to the volume of storage through the root
host computer, and wherein the at least one grandchild host computer has access to the
volume of storage through the at least one child host computer to create a logical
10 hierarchy of host computers that share access to the volume of storage.

28. The method of claim 27, further comprising acts of:

storing a first local copy of the at least a portion of the volume of storage
for the at least one child host computer; and
15 storing a second local copy of the at least a portion of the volume of
storage for the at least one grandchild host computer.

29. The method of claim 27, wherein the computer system has a physical
configuration, and wherein the logical hierarchy is independent of the physical
20 configuration of the computer system.

30. The method of claim 25, further comprising an act of caching a local copy
of the at least a portion of the volume of storage on the child host computer.

25 31. The method of claim 27, further comprising acts of:

caching a first local copy of the at least a portion of the volume of storage
on the at least one child host computer; and
caching a second local copy of the at least a portion of the volume of
storage on the at least one grandchild host computer.

32. The method of claim 25, wherein the at least one child host computer comprises a plurality of child host computers, and wherein at least one of the plurality of child host computers is untrusted by the root host computer.

5 33. — The method of claim 27, wherein the at least one child host computer has access to the volume of storage through the root host computer and the at least one grandchild host computer has access to the volume of storage through the at least one child host computer to create a logical hierarchy of host computers that share access to the shared volume, the logical hierarchy have at least a root level, a child level and a
10 grandchild level, and wherein the method further includes an act of:

(D) exporting the at least a portion of the volume of storage to at least one of the plurality of host computers in the logical hierarchy that is untrusted by the root host computer.

15 34. The method of claim 30, wherein the act of caching a local copy of the at least a portion of the volume of storage on the child host computer includes caching the at least a portion of the volume of storage in a set associative cache having at least two separately indexed portions.

20 35. The method of claim 25, wherein the volume of storage is a logical volume made available from the at least one storage system to the root host computer.

36. The method of claim 35, wherein the storage system has at least one disk drive that stores the volume of storage.

25 37. The method of claim 25, wherein the volume of storage is a logical volume made available from the storage system to the root host computer, and wherein the at least a portion of the volume of storage is a block of the logical volume.

30 38. The method of claim 27, wherein the computer system further comprises a network that interconnects the root host computer, the at least one child host computer and the at least one grandchild host computer.

39. The method of claim 27, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, and wherein the method further comprises an act of:

5 storing a local copy of at least portions of the volume of storage for each of the plurality of child host computers and the plurality of grandchild host computers.

40. The method of claim 27, further comprising an act of:

10 in response to a write request from one of the plurality of child host computers and the plurality of grandchild host computers to write to a first portion of the volume of storage, invalidating any local copies of the first portion of the volume of storage associated with the other child and grandchild host computers prior to authorizing the write request to proceed.

15

41. The method of claim 27, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy for distributing the shared volume of storage throughout the computer system, and wherein the method further comprises an act of:

20 providing each node in the logical hierarchy with information that enables it to participate in the hierarchy, wherein the information is limited to information relating only to adjacent nodes in the hierarchy.

25

42. The method of claim 25, further comprising an act of storing a local copy of the at least a portion of the volume of storage for the at least one child host computer.

30

43. The method of claim 27, wherein the at least one child host computer comprises a plurality of child host computers, wherein the at least one grandchild host computer comprises a plurality of grandchild host computers, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy

for distributing the shared volume of storage throughout the computer system, and wherein the method further comprises acts of:

storing a local copy of the volume of storage for each of the nodes in the hierarchy, each local copy comprising a copy of at least a portion of the volume of storage, wherein the at least a portion of the volume may differ between the nodes in the hierarchy; and

storing, in association with each local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the associated node to the at least a portion of the volume of storage stored in the local copy.

44. The method of claim 43, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for at least one second node in the hierarchy that is descendant from the at least one first node in the hierarchy.

45. The method of claim 43, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for each of the nodes in the hierarchy that is descendant from the at least one first node.

46. The method of claim 43, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the method further comprises acts of:

when a node in the hierarchy fails, recovering from the failure by reconnecting direct descendant nodes of the failed node to the parent node of the failed node; and

reestablishing the permission information for each of the direct descendant nodes of the failed node from the parent node of the failed node.

47. A method for use in a computer system including a plurality of host computers including a root host computer, at least one child host computer and at least

one grandchild host computer, the root host computer having at least one volume of storage available to it, the method comprising acts of:

(A) exporting at least a first portion of the volume of storage from the root host computer to the at least one child host computer; and

5 (B) exporting at least a second portion of the volume of storage from the child host computer to the at least one grandchild host computer, so that the at least one child host computer, the at least one grandchild host computer and the root host computer share access to the volume of storage.

10 48. The method of claim 47, wherein the act (A) includes an act of exporting the at least a first portion of the volume of storage to a plurality of child host computers so that the plurality of child host computers, the at least one grandchild host computer and the root host computer share access to the volume of storage.

15 49. The method of claim 48, wherein the act (B) includes an act of exporting the at least a second portion of the volume of storage to a plurality of grandchild host computers so that the plurality of child host computers, the plurality of grandchild host computers and the root host computer share access to the volume of storage.

20 50. The method of claim 47, wherein the act (B) includes an act of exporting the at least a second portion of the volume of storage to a plurality of grandchild host computers so that the at least one child host computer, the plurality of grandchild host computers and the root host computer share access to the volume of storage.

25 51. The method of claim 47, wherein the at least a second portion of the volume of storage is identical to the at least a first portion of the volume of storage.

52. The method of claim 47, wherein the at least a second portion of the volume of storage is a subset of the at least a first portion of the volume of storage.

30

53. The method of claim 47, further comprising an act of:

providing the at least one child host computer with access to the volume of storage only through the root host computer, and providing the at least one grandchild host computer with access to the volume of storage only through the at least one child host computer to create a logical hierarchy of host computers that share access to the
5 volume of storage.

54. The method of claim 47, wherein the computer system further comprises at least one storage system that makes the volume of storage available to the root host computer.

10

55. The method of claim 47, further comprising acts of:
caching a first local copy of the at least a first portion of the volume of storage on the at least one child host computer; and
caching a second local copy of the at least a second portion of the volume of storage on the at least one grandchild host computer.
15

56. The method of claim 53, further comprising an act of:
exporting at least a portion of the volume of storage to at least one of the plurality of host computers in the logical hierarchy that is untrusted by the root host computer.

20

57. The method of claim 55, wherein the act of caching a first local copy of the at least a first portion of the volume of storage on the at least one child host computer comprises caching the at least a first portion of the volume of in a set associative cache having at least two separately indexed portions; and
25 wherein the act of caching a second local copy of the at least a second portion of the volume of storage on the at least one grandchild host computer comprises caching the at least a second portion of the volume of in a set associative cache having two separately indexed portions.

30

58. The method of claim 54, wherein the volume of storage is a logical volume made available from the at least one storage system to the root host computer.

59. The method of claim 54, wherein the at least one storage system stores the
at least one volume of storage on a disk drive.

60. The method of claim 54, wherein the volume of storage is a logical
5 volume made available from the storage system to the root host computer, and wherein
the at least a first portion of the volume of storage is a block of the logical volume.

61. The method of claim 47, wherein the computer system further comprises
a network that interconnects the root host computer, the at least one child host computer
10 and the at least one grandchild host computer.

62. The method of claim 47, further comprising acts of:
storing a first local copy of the at least a first portion of the volume of
storage for the at least one child host computer; and
15 storing a second local copy of the at least a second portion of the volume
of storage for the at least one grandchild host computer.

63. The method of claim 62, further comprising an act of:
in response to a write request from one of the plurality of child host
20 computers and the plurality of grandchild host computers to write to a specified portion
of the volume of storage, invalidating any local copies of the specified portion of the
volume of storage associated with the other child and grandchild host computers prior to
authorizing the write request to proceed.

25 64. The method of claim 47, wherein each of the plurality of child and
grandchild host computers comprises a node in a logical hierarchy for distributing the
shared volume of storage throughout the computer system, and wherein the method
further comprises an act of:

30 providing each node in the logical hierarchy with information that enables
it to participate in the hierarchy, wherein the information is limited to information
relating only to adjacent nodes in the hierarchy.

65. The method of claim 47, further comprising an act of storing a local copy of the at least a first portion of the volume of storage for the at least one child host computer.

5 66. The method of claim 47, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy for distributing the shared volume of storage throughout the computer system, wherein the computer system has a physical configuration, and wherein the logical hierarchy is independent of the physical configuration of the computer system.

10

67. The method of claim 47, wherein each of the plurality of child and grandchild host computers comprises a node in a logical hierarchy for distributing the shared volume of storage throughout the computer system, and wherein the method further comprises acts of:

15 storing a local copy of the volume of storage for each of the nodes in the hierarchy, each local copy comprising a copy of at least a portion of the volume of storage, wherein the at least a portion of the volume may differ between the nodes in the hierarchy; and

20 storing, in association with each local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the associated node to the at least a portion of the volume of storage stored in the local copy.

25 68. The method of claim 67, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for at least one second node in the hierarchy that is descendant from the at least one first node in the hierarchy.

30 69. The method of claim 67, further comprising an act of storing, for at least one first node in the hierarchy, permissions information defining access authorization for each of the nodes in the hierarchy that is descendant from the at least one first node.

70. The method of claim 67, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the method further comprises acts of:

- when a node in the hierarchy fails, recovering from the failure by
- 5 reconnecting direct descendant nodes of the failed node to the parent node of the failed node; and
- reestablishing the permission information for each of the direct descendant nodes of the failed node from the parent node of the failed node.

10 71 A method for use in a computer system including a plurality of host computers including at least first and second root host computers, a first group of child host computers and a second group of child host computers, the first and second groups of child host computers each comprising at least one child host computer, the first and second root host computers each having a shared volume of storage available to it, the
15 method comprising acts of:

- (A) exporting at least a first portion of the shared volume of storage from the first root host computer to the first group of child host computers; and
- (B) exporting at least a second portion of the shared volume of storage from the second root host computer to the second group of child host computers, so that
20 the first and second root host computers and the first and second groups of child host computers all share access to the shared volume of storage.

72. The method of claim 71, further including an act of:

(C) in the event of a failure of the first root host computer, exporting
25 the at least a first portion of the shared volume from the second root host computer to the first group of child host computers.

73. The method of claim 71, further comprising an act of:

(D) making the shared volume available to the first and second root
30 host computers from a single copy of the shared volume.

74. The method of claim 71, wherein the computer system includes first and second storage systems, the first storage system storing a first copy of the shared volume and the second storage system storing a second copy of the shared volume, and wherein the method further comprises acts of:

5 (D) — making the first copy of the shared volume available to the first root host computer; and

(E) — making the second copy of the shared volume available to the second root host computer.

10 75. The method of claim 71, wherein the first and second portions of the shared volume are identical.

15 76. A computer readable medium encoded with a program for execution on a computer system including a plurality of host computers including a first host computer and at least one second host computer, the first host computer having a volume of storage available to it that is stored on at least one non-volatile storage device, the program, when executed, performs a method comprising an act of:

(A) — exporting at least a portion of the volume of storage from the first host computer to the at least one second host computer so that the at least one second host computer and the first host computer share access to the volume of storage.

20 77. The computer readable medium of claim 76, wherein the act (A) includes an act of exporting the at least a portion of the volume of storage to a plurality of second host computers so that the plurality of second host computers and the first host computer share access to the volume of storage.

78. The computer readable medium of claim 76, wherein the method further comprises an act of:

30 (B) — receiving the at least a portion of the volume of storage at the first host computer from a third host computer that exports the at least a portion of the volume of storage to the first host computer so that the at least one second host computer, the third host computer and the first host computer share access to the volume of storage,

wherein the at least one second host computer has access to the volume of storage through the first host computer, and wherein the first host computer has access to the volume of storage through the third host computer to create a logical hierarchy of host computers that share access to the volume of storage.

5

79. The computer readable medium of claim 76, wherein the at least one non-volatile storage device is provided on the first host computer.

80. The computer readable medium of claim 76, wherein the computer system
10 further comprises at least one storage system that includes the non-volatile storage
device and makes the volume of storage available to the first host computer.

81. The computer readable medium of claim 76, wherein the method further
comprises an act of caching a local copy of the at least a portion of the volume of storage
15 on the first host computer.

82. The computer readable medium of claim 76, wherein the at least one
second host computer comprises a plurality of second host computers, and wherein at
least one of the plurality of second host computers is untrusted by the first host
20 computer.

83. The computer readable medium of claim 81, wherein the act of caching a
local copy of the at least a portion of the volume of storage on the second host computer
includes caching the at least a portion of the volume of storage in a set associative cache
25 having two separately indexed portions.

84. The computer readable medium of claim 80, wherein the volume of
storage is a logical volume made available from the at least one storage system to the
first host computer.

30

85. The computer readable medium of claim 84, wherein the at least a portion
of the volume of storage is a block of the logical volume.

86. The computer readable medium of claim 76, wherein the method further comprises acts of:

storing a local copy of the at least a portion of the volume of storage for the first host computer; and

in response to a write request from one of the plurality of second host computers to write to a first portion of the volume of storage that is included within the at least a portion of the volume of storage stored within the local copy, invalidating the local copy of the first portion of the volume of storage.

10

87. The computer readable medium of claim 78, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers comprises a node in the logical hierarchy, wherein the computer system further includes additional host computers that are not directly coupled to the first host computer and form additional nodes in the logical hierarchy, and wherein the method further comprises an act of:

storing information for first host computer that enables it to participate in the hierarchy, wherein the information is limited to information relating only to the third host computer and the plurality of second host computers that are adjacent to the first host computer in the logical hierarchy.

88. The computer readable medium of claim 76, wherein the method further comprises an act of storing a local copy of the at least a portion of the volume of storage for the first host computer.

89. The computer readable medium of claim 78, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers comprises a node in the logical hierarchy, and wherein the method further comprises acts of:

storing a local copy of the volume of storage for the first host computer;
and

storing, in association with the local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the first host computer to the at least a portion of the volume of storage stored in the local copy.

5

90. The computer readable medium of claim 89, wherein the method further comprises an act of storing, for the first host computer, permissions information defining access authorization for at least one of the plurality of second host computers which are descendant from the first host computer in the hierarchy.

10

91. The computer readable medium of claim 89, wherein the method further comprises an act of storing, for the first host computer, permissions information defining access authorization for each of the plurality of second computers which are descendant from the first host computer in the hierarchy.

15

92. The computer readable medium of claim 89, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the method further comprises acts of:

20

when the third host computer fails, recovering from the failure by reconnecting the first host computer to the parent node of the third host computer; and reestablishing the permission information for the first host computer from the parent node of the failed node.

25

93. A computer readable medium encoded with a program for execution on a computer system including a plurality of host computers and at least one storage system, the plurality of host computers including a first host computer and at least one second host computer, the at least one storage system making a volume of storage available to the first host computer, the at least one storage system having at least storage device on which the volume of storage is stored, the program, when executed, performs a method comprising an act of:

30

(A) exporting at least a portion of the volume of storage from the first host computer to the at least one second host computer so that the at least one second host computer and the first host computer share access to the volume of storage.

5 94. The computer readable medium of claim 93, wherein the act (A) includes an act of exporting the at least a portion of the volume of storage to a plurality of second host computers so that the plurality of second host computers and the first host computer share access to the volume of storage.

10 95. The computer readable medium of claim 93, wherein the method further comprises an act of caching a local copy of the at least a portion of the volume of storage on the first host computer.

15 96. The computer readable medium of claim 93, wherein the at least one second host computer comprises a plurality of second host computers, and wherein at least one of the plurality of second host computers is untrusted by the first host computer.

20 97. The computer readable medium of claim 95, wherein the act of caching a local copy of the at least a portion of the volume of storage on the second host computer includes caching the at least a portion of the volume of storage in a set associative cache having at least two separately indexed portions.

25 98. The computer readable medium of claim 93, wherein the volume of storage is a logical volume made available from the at least one storage system to the first host computer.

30 99. The computer readable medium of claim 98, wherein the at least a portion of the volume of storage is a block of the logical volume.

100. The computer readable medium of claim 93, wherein the method further comprises acts of:

storing a local copy of the at least a portion of the volume of storage for the first host computer; and

- in response to a write request from one of the plurality of second host computers to write to a first portion of the volume of storage that is included within the at least a portion of the volume of storage stored within the local copy, invalidating the local copy of the first portion of the volume of storage.
- 5 portion of the volume of storage stored within the local copy, invalidating the local copy of the first portion of the volume of storage.

101. The computer readable medium of claim 93, wherein the at least one second host computer comprises a plurality of second host computers, wherein the plurality of second host computers have access to the volume of storage through the first host computer to create a logical hierarchy of nodes of host computers that share access to the volume of storage, wherein the computer system includes additional host computers that are not directly coupled to the first host computer and form additional nodes in the logical hierarchy, and wherein the method further comprises an act of:
- 15 storing information for the first host computer that enables it to participate in the hierarchy, wherein the information is limited to information relating only to the plurality of second host computers that are adjacent to the first host computer in the logical hierarchy.

- 20 102. The computer readable medium of claim 93, wherein the method further comprises an act of storing a local copy of the at least a portion of the volume of storage for the first host computer.

103. The computer readable medium of claim 93, wherein the at least one second host computer comprises a plurality of second host computers, wherein the plurality of second host computers have access to the volume of storage through the first host computer to create a logical hierarchy of nodes of host computers that share access to the volume of storage, and wherein the method further comprises acts of:
- 25 storing a local copy of the volume of storage for the first host computer;
- 30 and
- storing, in association with the local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions

information defining access authorization for the first host computer to at least a portion of the volume of storage stored in the local copy.

104. The computer readable medium of claim 103, wherein the method further
5 comprises an act of storing, for the first host computer, permissions information defining access authorization for at least one of the plurality of second host computers which are descendant from the first host computer in the logical hierarchy.

105. The computer readable medium of claim 103, wherein the method further
comprises an act of storing, for the first host computer, permissions information defining access authorization for each of the plurality of second host computers which are descendant from the first host computer in the logical hierarchy.

106. The computer readable medium of claim 103, wherein each of the nodes
15 in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, wherein the computer system includes at least one third host computer that is descendant from one of the plurality of second host computers in the logical hierarchy, and wherein the method further comprises acts of:

20 when the one of the second host computers fails, recovering from the failure by reconnecting the at least one third host computer to the first host computer which is the parent node of the failed one of the second host computers; and
reestablishing permission information for the third host computer from the first host computer.

107. A computer readable medium encoded with a program for execution on a computer system including a plurality of host computers including a first host computer, at least one second host computer, and a third host computer, the third host computer having a volume of storage available to it, the program, when executed, performs a
5 method comprising acts of:

(A) receiving at least a first portion of the volume of storage at the first host computer from the third host computer which exported the at least a first portion of the volume of storage to the first host computer so that the third host computer and the first host computer share access to the volume of storage; and

10 (B) exporting at least a second portion of the volume of storage from the first host computer to the at least one second host computer so that the at least one second host computer, the third host computer and the first host computer share access to the volume of storage.

15 108. The computer readable medium of claim 107, wherein the act (A) includes an act of receiving access for the first host computer to the volume of storage through the third host computer; and

wherein the act (B) includes an act of granting the at least one second host computer access to the volume of storage through the first host computer, whereby the at
20 least one second host computer the third host computer and the first host computer create a logical hierarchy of host computers that share access to the volume of storage.

109. The computer readable medium of claim 107, wherein the act (B) includes an act of exporting the at least a portion of the volume of storage to a plurality
25 of second host computers so that the plurality of second host computers, the first host computer and the third host computer share access to the volume of storage.

110. The computer readable medium of claim 107, wherein the method further comprises an act of caching a local copy of the at least a first portion of the volume of
30 storage on the first host computer.

111. The computer readable medium of claim 107, wherein the at least one second host computer comprises a plurality of second host computers, and wherein at least one of the plurality of second host computers is untrusted by the first host computer.

5

112. The computer readable medium of claim 110, wherein the act of caching a local copy of the at least a first portion of the volume of storage on the second host computer includes caching the at least a first portion of the volume of storage in a set associative cache having at least two separately indexed portions.

10

113. The computer readable medium of claim 107, wherein the method further comprises acts of:

storing a local copy of the at least a first portion of the volume of storage for the first host computer; and

15

in response to a write request from one of the plurality of second host computers to write to a specified portion of the volume of storage that is included within the at least a first portion of the volume of storage stored in the local copy, invalidating the local copy of the specified portion of the volume of storage.

20

114. The computer readable medium of claim 108, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers comprises a node in the logical hierarchy, wherein the computer system further includes additional host computers that are not directly coupled to the first host computer and form

25

additional nodes in the logical hierarchy, and wherein the method further comprises an act of:

storing information for first host computer that enables it to participate in the logical hierarchy, wherein the information is limited to information relating only to the third host computer and the plurality of second host computers that are adjacent to

30

the first host computer in the logical hierarchy.

115. The computer readable medium of claim 107, wherein the method further comprises an act of storing a local copy of the at least a first portion of the volume of storage for the first host computer.

5 116. The computer readable medium of claim 108, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers each comprises a node in the logical hierarchy, and wherein the method further comprises acts of:

10 storing a local copy of the at least a first portion of the volume of storage for the first host computer; and

 storing, in association with the local copy, metadata describing the at least a first portion of the volume of storage included in the local copy and permissions information defining access authorization for the first host computer to the at least a first portion of the volume of storage included in the local copy.

15 117. The computer readable medium of claim 116, wherein the method further comprises an act of storing, for the first host computer, permissions information defining access authorization for at least one of the plurality of second host computers which are descendant from the first host computer in the logical hierarchy.

20 118. The computer readable medium of claim 116, wherein the method further comprises an act of storing, for the first host computer, permissions information defining access authorization for each of the plurality of second host computers which are descendant from the first host computer in the logical hierarchy.

25 119. The computer readable medium of claim 116, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the method further comprises acts of:

 when the third host computer fails, recovering from the failure by reconnecting the first host computer to the parent node of the third host computer; and

reestablishing the permission information for the first host computer from the parent node of the third host computer.

120. The computer readable medium of claim 107, wherein the at least a
5 second portion of the volume of storage is identical to the at least a first portion of the
volume of storage.

121. The computer readable medium of claim 107, wherein the at least a
second portion of the volume of storage is a subset of the at least a first portion of the
10 volume of storage.

122. A first host computer for use in a computer system including a plurality of
host computers including the first host computer and at least one second host computer,
the first host computer having a volume of storage available to it that is stored on at least
15 one non-volatile storage device, the first host computer comprising:

at least one port that enables the first host computer to be coupled to other
components in the computer system; and

20 at least one controller, coupled to the at least one port, to export at least a portion
of the volume of storage from the first host computer to the at least one second host
computer so that the at least one second host computer and the first host computer can
share access to the volume of storage.

123. The first host computer of claim 122, wherein the at least one controller
comprises means for exporting the at least a portion of the volume of storage from the
25 first host computer to the at least one second host computer so that the at least one
second host computer and the first host computer can share access to the volume of
storage

124. The first host computer of claim 122, wherein the at least one controller
30 exports the at least a portion of the volume of storage to a plurality of second host
computers so that the plurality of second host computers and the first host computer can
share access to the volume of storage.

125. The first host computer of claim 122, wherein the at least one controller receives the at least a portion of the volume of storage from a third host computer that exports the at least a portion of the volume of storage to the first host computer so that
5 the third host computer and the first host computer share access to the volume of storage, wherein the at least one controller grants access to the at least a portion of the volume of storage to the at least one second host computer through the first host computer, and wherein the first host computer receives access to the volume of storage through the third host computer, whereby a logical hierarchy of host computers is created that share access
10 to the volume of storage.

126. The first host computer of claim 122, further comprising the at least one non-volatile storage device.

15 127. The first host computer of claim 122, wherein the computer system further comprises at least one storage system that includes the non-volatile storage device and makes the volume of storage available to the first host computer, and wherein the at least one controller is adapted to receive the volume of storage from the at least one storage system.
20

128. The first host computer of claim 122, further comprising a cache, and wherein the at least one controller stores a local copy of the at least a portion of the volume of storage in the cache.

25 129. The first host computer of claim 122, wherein the at least one second host computer comprises a plurality of second host computers, and wherein the at least one controller exports at least a portion of the volume of storage to at least one of the plurality of second host computers that is untrusted by the first host computer.

30 130. The first host computer of claim 128, wherein the cache is a set associative cache having at least two separately indexed portions.

131. The first host computer of claim 127, wherein the volume of storage is a logical volume made available from the at least one storage system to the first host computer, and wherein the at least one controller is adapted to receive the logical volume from the at least one storage system.

5

132. The first host computer of claim 122, wherein the at least one controller stores a local copy of the at least a portion of the volume of storage for the first host computer; and

10 wherein in response to a write request from one of the plurality of second host computers to write to a first portion of the volume of storage that is included within the at least a portion of the volume of storage stored within the local copy, the at least one controller invalidates the local copy of the first portion of the volume of storage.

15 133. The first host computer of claim 125, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers comprises a node in the logical hierarchy, wherein the computer system further includes additional host computers that are not directly coupled to the first host computer and form additional nodes in the logical hierarchy, and wherein the at least one controller stores information 20 for first host computer that enables it to participate in the hierarchy, the information being limited to information relating only to the third host computer and the plurality of second host computers that are adjacent to the first host computer in the logical hierarchy.

25 134. The first host computer of claim 122, wherein the at least one controller stores a local copy of the at least a portion of the volume of storage for the first host computer.

30 135. The first host computer of claim 125, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers each comprises a node in

the logical hierarchy, wherein the at least one controller stores a local copy of the volume of storage for the first host computer; and

- wherein the at least one controller stores, in association with the local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the first host computer to the at least a portion of the volume of storage stored in the local copy.
- 5

136. The first host computer of claim 135, wherein the at least one controller stores, for the first host computer, permissions information defining access authorization 10 for at least one of the plurality of second host computers which are descendant from the first host computer in the hierarchy.

137. The first host computer of claim 135, wherein the at least one controller stores, for the first host computer, permissions information defining access authorization 15 for each of the plurality of second computers which are descendant from the first host computer in the hierarchy.

138. The first host computer of claim 135, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a 20 parent node to its descendant nodes in the hierarchy, and wherein the at least one controller, when the third host computer fails, recovers from the failure by reconnecting the first host computer to the parent node of the third host computer and reestablishing the permission information for the first host computer from the parent node of the third host computer.

25

139. A first host computer for use in a computer system including a plurality of host computers and at least one storage system, the plurality of host computers including the first host computer and at least one second host computer, the at least one storage system making a volume of storage available to the first host computer, the at least one storage system having at least storage device on which the volume of storage is stored, 30 the first host computer comprising:

at least one port that enables the first host computer to be coupled to other components in the computer system; and

at least one controller, coupled to the at least one port, to export at least a portion of the volume of storage from the first host computer to the at least one second host

- 5 computer so that the at least one second host computer and the first host computer can share access to the volume of storage.

140. The first host computer of claim 139, wherein the at least one controller comprises means for exporting the at least a portion of the volume of storage from the
10 first host computer to the at least one second host computer so that the at least one second host computer and the first host computer can share access to the volume of storage

141. The first host computer of claim 139, wherein the at least one controller exports the at least a portion of the volume of storage to a plurality of second host computers so that the plurality of second host computers and the first host computer can share access to the volume of storage.

142. The first host computer of claim 139, further comprising a cache, and
20 wherein the at least one controller stores a local copy of the at least a portion of the volume of storage in the cache.

143. The first host computer of claim 139, wherein the at least one second host computer comprises a plurality of second host computers, and wherein the at least one
25 controller exports at least a portion of the volume of storage to at least one of the plurality of second host computers that is untrusted by the first host computer.

144. The first host computer of claim 142, wherein the cache is a set associative cache having at least two separately indexed portions.

30

145. The first host computer of claim 139, wherein the volume of storage is a logical volume made available from the at least one storage system to the first host

computer, and wherein the at least one controller is adapted to receive the logical volume from the at least one storage system.

146. The first host computer of claim 139, wherein the at least one controller
5 stores a local copy of the at least a portion of the volume of storage for the first host computer; and

wherein in response to a write request from one of the plurality of second host computers to write to a first portion of the volume of storage that is included within the at least a portion of the volume of storage stored within the local copy, the at least one 10 controller invalidates the local copy of the first portion of the volume of storage.

147. The first host computer of claim 139, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first host computer each comprises a node in the 15 logical hierarchy, wherein the computer system further includes additional host computers that are not directly coupled to the first host computer and form additional nodes in the logical hierarchy, and wherein the at least one controller stores information for first host computer that enables it to participate in the hierarchy, the information being limited to information relating only to the plurality of second host computers that 20 are adjacent to the first host computer in the logical hierarchy.

148. The first host computer of claim 139, wherein the at least one controller stores a local copy of the at least a portion of the volume of storage for the first host computer.

25
149. The first host computer of claim 139, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first host computer each comprises a node in the logical hierarchy, wherein the at least one controller stores a local copy of the volume of 30 storage for the first host computer; and

wherein the at least one controller stores, in association with the local copy, metadata describing the at least a portion of the volume of storage included in the

local copy and permissions information defining access authorization for the first host computer to the at least a portion of the volume of storage stored in the local copy.

150. The first host computer of claim 149, wherein the at least one controller
5 stores, for the first host computer, permissions information defining access authorization for at least one of the plurality of second host computers which is descendant from the first host computer in the hierarchy.

151. The first host computer of claim 149, wherein the at least one controller
10 stores, for the first host computer, permissions information defining access authorization for each of the plurality of second computers which are descendant from the first host computer in the hierarchy.

152. The first host computer of claim 149, wherein each of the nodes in the
15 logical hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, wherein the computer system includes at least one third host computer that is descendant from one of the plurality of second host computers in the logical hierarchy, and wherein when the one of the second host computers fails, the at least one controller recovers from the failure by reconnecting
20 the at least one third host computer to the first host computer and reestablishing permission information for the third host computer from the first host computer.

153. A first host computer for use in a computer system including a plurality of host computers including the first host computer, at least one second host computer, and
25 a third host computer, the third host computer having a volume of storage available to it, the first host computer comprising:

at least one port that enables the first host computer to be coupled to other components in the computer system; and

30 at least one controller, coupled to the at least one port, to receive at least a first portion of the volume of storage from the third host computer which exports the at least a first portion of the volume of storage to the first host computer so that the third host computer and the first host computer can share access to the volume of storage, the at

least one controller further adapted to export at least a second portion of the volume of storage from the first host computer to the at least one second host computer so that the at least one second host computer, the third host computer and the first host computer can share access to the volume of storage.

5

154. The first host computer of claim 153, wherein the at least one controller comprises means for receiving the at least a first portion of the volume of storage from the third host computer and means for exporting the at least a second portion of the volume of storage from the first host computer to the at least one second host computer
10 so that the at least one second host computer, the third host computer and the first host computer can share access to the volume of storage

155. The first host computer of claim 153, wherein the at least one controller is adapted to receive access to the volume of storage through the third host computer and
15 to grant the at least one second host computer access to the volume of storage through the first host computer, whereby the at least one second host computer, the third host computer and the first host computer can form a logical hierarchy of host computers that share access to the volume of storage.

20

156. The first host computer of claim 153, wherein the at least one controller is adapted to export the at least a portion of the volume of storage to a plurality of second host computers so that the plurality of second host computers, the first host computer and the third host computer can share access to the volume of storage.

25

157. The first host computer of claim 153, further comprising a cache, and wherein the at least one controller is adapted to store a local copy of the at least a portion of the volume of storage in the cache.

30

158. The first host computer of claim 153, wherein the at least one second host computer comprises a plurality of second host computers, and wherein the at least one controller is adapted to export at least a portion of the volume of storage to at least one of the plurality of second host computers that is untrusted by the first host computer.

159. The first host computer of claim 157, wherein the cache is a set associative cache having at least two separately indexed portions.

5 160. The first host computer of claim 153, wherein the volume of storage is a logical volume made available from the at least one storage system to the first host computer, and wherein the at least one controller is adapted to receive the logical volume from the at least one storage system.

10 161. The first host computer of claim 153, wherein the at least one controller is adapted to store a local copy of the at least a portion of the volume of storage for the first host computer; and

15 wherein in response to a write request from one of the plurality of second host computers to write to a first portion of the volume of storage that is included within the at least a portion of the volume of storage stored within the local copy, the at least one controller is adapted to invalidate the local copy of the first portion of the volume of storage.

162. The first host computer of claim 153, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers each comprises a node in the logical hierarchy, wherein the computer system further includes additional host computers that are not directly coupled to the first host computer and form additional nodes in the logical hierarchy, and wherein the at least one controller stores information for first host computer that enables it to participate in the hierarchy, the information being limited to information relating only to the third host computer and the plurality of second host computers that are adjacent to the first host computer in the logical hierarchy.

30 163. The first host computer of claim 153, wherein the at least one controller is adapted to store a local copy of the at least a portion of the volume of storage for the first host computer.

164. The first host computer of claim 153, wherein the at least one second host computer comprises a plurality of second host computers, wherein each of the plurality of second host computers and the first and third host computers each comprises a node in the logical hierarchy, wherein the at least one controller is adapted to store a local copy
5 of the volume of storage for the first host computer; and

wherein the at least one controller is adapted to store, in association with the local copy, metadata describing the at least a portion of the volume of storage included in the local copy and permissions information defining access authorization for the first host computer to the at least a portion of the volume of storage stored in the
10 local copy.

165. The first host computer of claim 164, wherein the at least one controller is adapted to store, for the first host computer, permissions information defining access authorization for at least one of the plurality of second host computers which are
15 descendant from the first host computer in the hierarchy.

166. The first host computer of claim 164, wherein the at least one controller is adapted to store, for the first host computer, permissions information defining access authorization for each of the plurality of second computers which are descendant from
20 the first host computer in the hierarchy.

167. The first host computer of claim 164, wherein each of the nodes in the hierarchy has a parent node, wherein the permissions information is granted from a parent node to its descendant nodes in the hierarchy, and wherein the at least one
25 controller, when the third host computer fails, is adapted to recover from the failure by reconnecting the first host computer to the parent node of the third host computer and reestablishing the permission information for the first host computer from the parent node of the third host computer.

30 168. The first host computer of claim 153, wherein the at least a second portion of the volume of storage is identical to the at least a first portion of the volume of storage.

169. The first host computer of claim 153, wherein the at least a second portion of the volume of storage is a subset of the at least a first portion of the volume of storage.